

DIGITAL FORENSICS: 5 WAYS TO SPOT A FAKE PHOTO

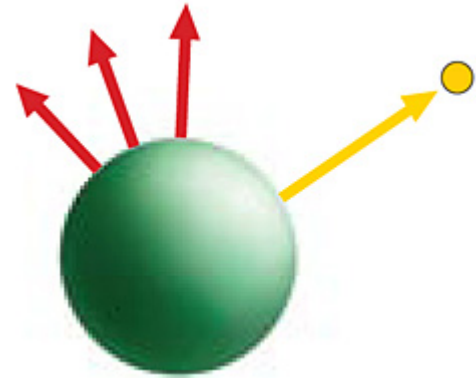
By Hany Farid

This story is a supplement to the feature "[Digital Forensics: How Experts Uncover Doctored Images](#)" which was printed in the [June 2008](#) issue of Scientific American.

Lighting

Composite images made of pieces from different photographs can display subtle differences in the lighting conditions under which each person or object was originally photographed. Such discrepancies will often go unnoticed by the naked eye.

For an image such as the one at the right, my group can estimate the direction of the light source for each person or object (arrows). Our method relies on the simple fact that the amount of light striking a surface depends on the relative orientation of the surface to the light source. A sphere, for example, is lit the most on the side facing the light and the least on the opposite side, with gradations of shading across its surface according to the angle between the surface and the direction to the light at each point.



To infer the light-source direction, you must know the local orientation of the surface. At most places on an object in an image, it is difficult to determine the orientation. The one exception is along a surface contour, where the orientation is perpendicular to the contour (*red arrows right*). By measuring the brightness and orientation along several points on a contour, our algorithm estimates the light-source direction.




DIGITAL FORENSICS: 5 WAYS TO SPOT A FAKE PHOTO

By Hany Farid

This story is a supplement to the feature "[Digital Forensics: How Experts Uncover Doctored Images](#)" which was printed in the [June 2008](#) issue of Scientific American.

For the image above, the light-source direction for the police does not match that for the ducks (arrows). We would have to analyze other items to be sure it was the ducks that were added.



Eyes and Positions

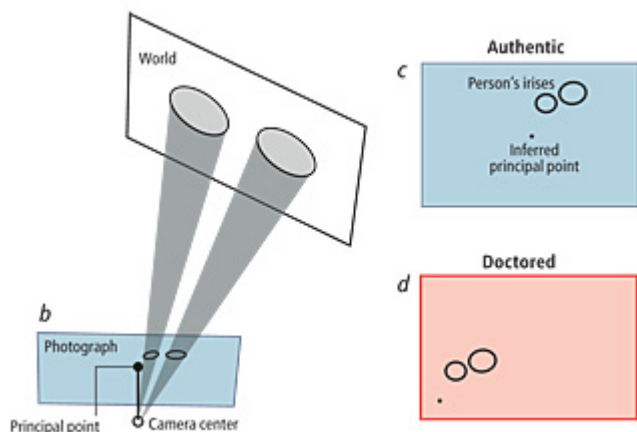
Because eyes have very consistent shapes, they can be useful for assessing whether a photograph has been altered.

A person's irises are circular in reality but will appear increasingly elliptical as the eyes turn to the side or up or down (a). One can approximate how eyes will look in a photograph by tracing rays of light running from them to a point called the camera center (b). The picture forms where the rays cross the image plane (blue). The principal point of the camera—the intersection of the image plane and the ray along which the camera is pointed—will be near the photograph's center.

My group uses the shape of a person's two irises in the photograph to infer how his or her eyes are

oriented relative to the camera and thus where the camera's principal point is located (c). A principal point far from the center or people having inconsistent principal points is evidence of tampering (d). The algorithm also works with other objects if their shapes are known, as with two wheels on a car.

The technique is limited, however, because the analysis relies on accurately measuring the slightly different shapes of a person's two irises. My collaborators and I have found we can reliably estimate large camera differences, such as when a person is moved from one side of the image to the middle. It is harder to tell if the person was moved much less than that.



Specular

Surrounding lights reflect in eyes to form small white dots called specular highlights. The shape, color and location of these highlights tell us quite a bit about the lighting.

Highlights

DIGITAL FORENSICS: 5 WAYS TO SPOT A FAKE PHOTO

By Hany Farid

This story is a supplement to the feature "[Digital Forensics: How Experts Uncover Doctored Images](#)" which was printed in the [June 2008](#) issue of Scientific American.




In 2006 a photo editor contacted me about a picture of *American Idol* stars that was scheduled for publication in his magazine (*above*). The specular highlights were quite different (*insets*).

DIGITAL FORENSICS: 5 WAYS TO SPOT A FAKE PHOTO


By Hany Farid

This story is a supplement to the feature "[Digital Forensics: How Experts Uncover Doctored Images](#)" which was printed in the [June 2008](#) issue of Scientific American.



The highlight position indicates where the light source is located (*above left*). As the direction to the light source (*yellow arrow*) moves from left to right, so do the specular highlights.

The highlights in the *American Idol* picture are so inconsistent that visual inspection is enough to infer the photograph has been doctored. Many cases, however, require a mathematical analysis. To determine light position precisely requires taking into account the shape of the eye and the relative orientation between the eye, camera and light. The orientation matters because eyes are not perfect spheres: the clear covering of the iris, or cornea, protrudes, which we model in software as a sphere whose center is offset from the center of the whites of the eye, or sclera (*above right*).



whose center is offset from the center of the whites of the eye, or sclera (*above right*).

Our algorithm calculates the orientation of a person's eyes from the shape of the irises in the image. With this information and the position of the specular highlights, the program estimates the direction to the light. The image of the *American Idol* cast (*above; directions depicted by red dots on green spheres*) was very likely composed from at least three photographs.

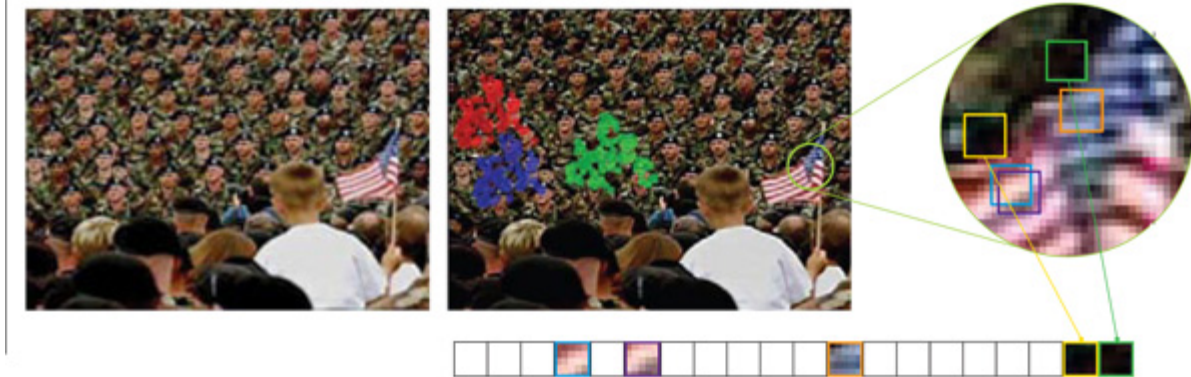
Send in the Clones

Cloning—the copying and pasting of a region of an image—is a very common and powerful form of manipulation.

DIGITAL FORENSICS: 5 WAYS TO SPOT A FAKE PHOTO

By Hany Farid

This story is a supplement to the feature "[Digital Forensics: How Experts Uncover Doctored Images](#)" which was printed in the [June 2008](#) issue of Scientific American.



This image is taken from a television ad used by George W. Bush's reelection campaign late in 2004. Finding cloned regions by a brute-force computer search, pixel by pixel, of all possible duplicated regions is impractical because they could be of any shape and located anywhere in the image. The number of comparisons to be made is astronomical, and innumerable tiny regions will be identical just by chance ("false positives"). My group has developed a more efficient technique that works with small blocks of pixels, typically about a six-by-six-pixel square (*inset*).

For every six-by-six block of pixels in the image, the algorithm computes a quantity that characterizes the colors of the 36 pixels in the block. It then uses that quantity to order all the blocks in a sequence that has identical and very similar blocks close together. Finally, the program looks for the identical blocks and tries to "grow" larger identical regions from them block by block. By dealing in blocks, the algorithm greatly reduces the number of false positives that must be examined and discarded.

When the algorithm is applied to the image from the political ad, it detects three identical regions (red, blue and green).

Camera Fingerprints

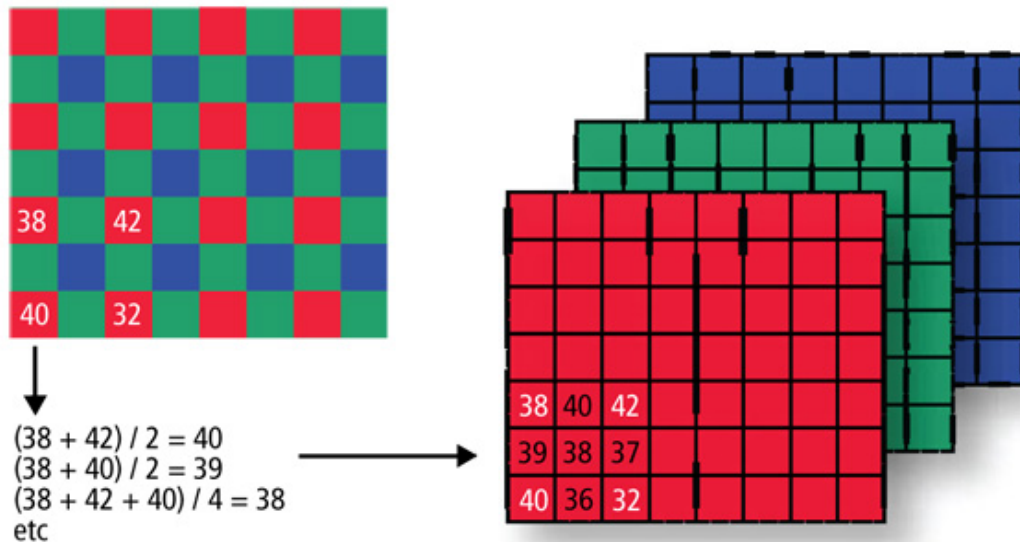
Digital retouching rarely leaves behind a visual trace. Because retouching can take many forms, I wanted to develop an algorithm that would detect any modification of an image. The technique my group came up with depends on a feature of how virtually all digital cameras work.

A camera's digital sensors are laid out in a rectangular grid of pixels, but each pixel detects the intensity of light only in a band of wavelengths near one color, thanks to a color filter array (CFA) that sits on top of the digital sensor grid. The CFA used most often, the Bayer array, has red, green and blue filters arranged as shown below.

DIGITAL FORENSICS: 5 WAYS TO SPOT A FAKE PHOTO

By Hany Farid

This story is a supplement to the feature "[Digital Forensics: How Experts Uncover Doctored Images](#)" which was printed in the [June 2008](#) issue of Scientific American.



Each pixel in the raw data thus has only one color channel of the three required to specify a pixel of a standard digital image. The missing data are filled in—either by a processor in the camera itself or by software that interprets raw data from the camera—by interpolating from the nearby pixels, a procedure called demosaicing. The simplest approach is to take the average of neighboring values, but more sophisticated algorithms are also used to achieve better results. Whatever demosaicing algorithm is applied, the pixels in the final digital image will be correlated with their neighbors. If an image does not have the proper pixel correlations for the camera allegedly used to take the picture, the image has been retouched in some fashion.



DIGITAL FORENSICS: 5 WAYS TO SPOT A FAKE PHOTO

By Hany Farid

This story is a supplement to the feature "[Digital Forensics: How Experts Uncover Doctored Images](#)" which was printed in the [June 2008](#) issue of Scientific American.

My group's algorithm looks for these periodic correlations in a digital image and can detect deviations from them. If the correlations are absent in a small region, most likely some spot changes have been made there. The correlations may be completely absent if image-wide changes were made, such as resizing or heavy JPEG compression. This technique can detect changes such as those made by Reuters to an image it released from a meeting of the United Nations Security Council in 2005 (*above*): the contrast of the notepad was adjusted to improve its readability.

A drawback of the technique is that it can be applied usefully only to an allegedly original digital image; a scan of a printout, for instance, would have new correlations imposed courtesy of the scanner.